

CIFTS Workflows

IU Contribution

Joshua Hursey, Abhishek Kulkarni,
Timothy I. Mattox, Torsten Hoefler, Andrew Lumsdaine
Open Systems Laboratory
Indiana University
Bloomington, IN 47405
{jjhursey,adkulkar,timattox,htor,lums}@osl.iu.edu

February 20, 2009

Summary of Workflows

We believe that many of the workflows in the proposal are still applicable, but with some slight modifications. Most of the workflow discussion in this document focuses on the role of MPI (particularly Open MPI). Further iterations are needed to refine these workflows such that they are correct for other components of the CIFTS FTB.

Workflow: Node Failure

All of these workflows detail a response to a detected node failure.

- Section 1 Details the registered events for various components.
- Section 1.1 Node failure without a job
- Section 1.2 Node failure with MPI job aborting
- Section 1.3 Node failure with MPI job continuing

Workflow: Checkpoint/Restart & Process Migration

All of these workflows detail a response to a predicted node failure. So with advance notice of a failure, preventative actions are triggered to mitigate the impact of the failure. Additionally a RM/JS might wish to trigger a checkpoint to provide a coarse-grained, gang scheduling type of functionality.

- Section 2 Details the registered events for various components.
- Section 2.1 Gang Scheduling Support
- Section 2.2 Predicted node failure, resulting in a full job suspension/shutdown
- Section 2.3 Predicted node failure, resulting in process migration

Workflow: Interconnect Failure

All of these workflows detail a response to a faulty interconnect.

- Section 3 Details the registered events for various components.
- Section 3.1 Fail-over to an alternative device.
- Section 3.2 React to corrupted or missing data

Workflow: Task Farm

The task farm workflow concerns an MPI application that operates in a manager/worker model. This workflow still needs to be more concretely specified in a later draft.

1 Workflow: Node Failure

The following table details the events that each component will want to either throw or catch.

	Component	Action	Message
	<i>Initialization & Job Launch</i>		
0	RM/JS	Register	Check Problem Node (node *)
0	RM/JS	Register	Dead Physical Node (node *)
0	RM/JS	Register	Dead MPI Node (node *: job z)
0	RM/JS	Register	Restored Node (node *)
0	RM/JS	Register	Restored MPI Node (node *: job z)
0	Monitoring System	Register	Check Problem Node (node *)
0	Monitoring System	Register	Dead Physical Node (node *)
0	Monitoring System	Register	Restored Node (node *)
0	Autonomic Script	Register	Check Problem Node (node *)
0	Autonomic Script	Register	Dead Physical Node (node *)
0	MPI	Register	Dead MPI Node (node *: job z)
0	MPI	Register	Restored MPI Node (node *: job z)
0	MPI	Register	Dead MPI Rank (node x: job z: rank n-m)
0	Application	Register	Dead MPI Rank (node x: job z: rank n-m)
0	Application	Register	Restored MPI Node (node x: job z)

1.1 Workflow: Node Failure Without Job

A node failure can occur without any jobs running on the failed node.

	Component	Action	Message
	<i>Node x Fails, no job running on node x</i>		
1	Monitoring System	Throw	Check Problem Node (node x)
	<i>Suspect problem with node x</i>		
2(a)	RM/JS	Catch	Check Problem Node (node x)
	<i>Suspend scheduling on node x (suspect failure)</i>		
2(b)	Autonomic Script	Catch	Check Problem Node (node x)
	<i>Attempt to confirm node x failed</i>		
3	Autonomic Script	Throw	Dead Physical Node (node x)
	<i>Confirmed node x failed</i>		
4(a)	RM/JS	Catch	Dead Physical Node (node x)
	<i>Remove node x from resource pool</i>		
4(b)	Monitoring System	Catch	Dead Physical Node (node x)
	<i>Remove node x from set of monitored resources</i>		
4(c)	Autonomic Script	Catch	Dead Physical Node (node x)
	<i>Notify sysadmin, trigger full diagnosis after hard reboot</i>		
	<i>Archives system logs, begin stress test, bring online spare nodes</i>		
	<i>Refund CPU accounting units, reschedule job</i>		
	<i>Time passes, machine returned to service</i>		
5	Autonomic Script	Throw	Restored Node (node x)
	<i>Sysadmin uses script to notify services of node recovery</i>		
6(a)	RM/JS	Catch	Restored Node (node x)
	<i>Return node x to resource pool</i>		
6(b)	Monitoring System	Catch	Restored Node (node x)
	<i>Return node x to the set of monitored resources</i>		

1.2 Workflow: Node Failure With MPI Job Aborting

A node failure occurs while a job is running on the failed node. The policy expressed by the application through the MPI interface is that the MPI abort on such a failure.

	Component	Action	Message
	<i>Node x Fails, job z running on allocation including node x</i>		
1	Monitoring System	Throw	Check Problem Node (node x)
	<i>Suspect problem with node x</i>		
2(a)	RM/JS	Catch	Check Problem Node (node x)
	<i>Mark node x as (suspect failure)</i>		
2(b)	Autonomic Script	Catch	Check Problem Node (node x)
	<i>Attempt to confirm node x failed</i>		
3	Autonomic Script	Throw	Dead Physical Node (node x)
	<i>Confirmed node x failed</i>		
4(a)	RM/JS	Catch	Dead Physical Node (node x)
	<i>Remove node x from resource pool</i>		
4(b)	Monitoring System	Catch	Dead Physical Node (node x)
	<i>Remove node x from set of monitored resources</i>		
4(c)	Autonomic Script	Catch	Dead Physical Node (node x)
	<i>Notify sysadmin, trigger full diagnosis after hard reboot</i>		
	<i>Archives system logs, begin stress test, bring online spare nodes</i>		
	<i>Refund CPU accounting units, reschedule job</i>		
5	RM/JS	Throw	Dead MPI Node (node x: job z)
	<i>Translates node x to job z</i>		
6	MPI	Catch	Dead MPI Node (node x: job z)
	<i>MPI prints console error, aborts job z</i>		
	<i>Time passes, machine returned to service</i>		
7	Autonomic Script	Throw	Restored Node (node x)
	<i>Sysadmin uses script to notify services of node recovery</i>		
8(a)	RM/JS	Catch	Restored Node (node x)
	<i>Return node x to unallocated resource pool</i>		
8(b)	Monitoring System	Catch	Restored Node (node x)
	<i>Return node x to the set of monitored resources</i>		

1.3 Workflow: Node Failure With MPI Job Continuing

A node failure occurs while a job is running on the failed node. Node failure policy is that MPI should continue with holes in communicators. Node recovery policy is that MPI adds resources to internal pool to support application directed re-spawning of processes.

	Component	Action	Message
	<i>Node x Fails, job z running on allocation including node x</i>		
1	Monitoring System	Throw	Check Problem Node (node x) <i>Suspect problem with node x</i>
2(a)	RM/JS	Catch	Check Problem Node (node x) <i>Mark node x as (suspect failure)</i>
2(b)	Autonomic Script	Catch	Check Problem Node (node x) <i>Attempt to confirm node x failed</i>
3	Autonomic Script	Throw	Dead Physical Node (node x) <i>Confirmed node x failed</i>
4(a)	RM/JS	Catch	Dead Physical Node (node x) <i>Remove node x from resource pool</i>
4(b)	Monitoring System	Catch	Dead Physical Node (node x) <i>Remove node x from set of monitored resources</i>
4(c)	Autonomic Script	Catch	Dead Physical Node (node x) <i>Notify sysadmin, trigger full diagnosis after hard reboot</i> <i>Archives system logs, begin stress test, bring online spare nodes</i> <i>Refund CPU accounting units, reschedule job</i>
5	RM/JS	Throw	Dead MPI Node (node x: job z) <i>Translates node x to job z</i>
6	MPI	Catch	Dead MPI Node (node x: job z) <i>Translate (node x:job z) to ranks m-n</i> <i>Replace ranks m-n with MPI_PROC_NULL, call application error handlers</i>
7	MPI	Throw	Dead MPI Rank (node x: job z: rank n-m) <i>Translate (node x:job z) to ranks m-n</i>
8	Application	Catch	Dead MPI Rank (node x: job z: rank n-m) <i>Work around 'blank' ranks n-m in the MPI communicators</i>
	<i>Time passes, machine returned to service</i>		
9	Autonomic Script	Throw	Restored Node (node x) <i>Sysadmin uses script to notify services of node recovery</i>
10(a)	RM/JS	Catch	Restored Node (node x) <i>Return node x to resource pool for job z</i>
10(b)	Monitoring System	Catch	Restored Node (node x) <i>Return node x to the set of monitored resources</i>
11	RM/JS	Throw	Restored MPI Node (node x: job z) <i>Translates node x to job z</i>
12(a)	MPI	Catch	Restored MPI Node (node x: job z) <i>Add node x as an unallocated resource</i>
12(b)	Application	Catch	Restored MPI Node (node x: job z) <i>If needed, use MPI_Comm_spawn to create new processes</i>

2 Workflow: Checkpoint/Restart & Process Migration

All of these workflows detail a response to a predicted node failure. So with advance notice of a failure, preventative actions are triggered to mitigate the impact of the failure. Additionally a RM/JS might wish to trigger a checkpoint to provide a coarse-grained, gang scheduling type of functionality.

	Component	Action	Message
	<i>Initialization & Job Launch</i>		
0	RM/JS	Register	Restored Node (node *)
0	RM/JS	Register	Suspend Job (job z)
0	RM/JS	Register	Resume Job (job z)
0	RM/JS	Register	Resume Job Cmd (job z)
0	RM/JS	Register	Predict Problem Node (node *)
0	RM/JS	Register	Migrate Node (job z: node x,q)
0	RM/JS	Register	Migrate Node Done (job z: node x,q)
0	Autonomic Script	Register	Restored Node (node *)
0	Autonomic Script	Register	Predict Problem Node (node *)
0	MPI	Register	Suspend Job (job z)
0	MPI	Register	Resume Job (job z)
0	MPI	Register	Resume Job Cmd (job z)
0	MPI	Register	Migrate Node (job z: node x,q)
0	MPI	Register	Migrate Node Done (job z: node x,q)

2.1 Workflow: Gang Scheduling Support

Gang scheduling support. The RM/JS suspends and resumes entire jobs using a checkpoint/restart technique in cooperation with the MPI implementation.

	Component	Action	Message
	<i>RM/JS decides to suspend job z using CPR</i>		
1	RM/JS	Throw	Suspend Job (job z)
	<i>Suspend job z</i>		
2	MPI	Catch	Suspend Job (job z)
	<i>Coordinate a global checkpoint operation. Suspend/Terminate job z</i>		
3	MPI	Throw	Resume Job Cmd (job z)
	<i>Provide RM/JS with the command needed to resume job z</i>		
4	RM/JS	Catch	Resume Job Cmd (job z)
	<i>Store command with information for job z</i>		
	<i>RM/JS decides to resume job z from CPR</i>		
5	RM/JS	Throw	Resume Job (job z)
	<i>Use stored resume information for job z to restart job</i>		
6	MPI	Catch	Resume Job (job z)
	<i>Bring job z back into a running state</i>		

2.2 Workflow: Predicted Failure, Job Suspend

A monitoring system predicts a node failure based on heuristic information gathered from the operating system, network card, and other system resources. The job is suspended and rescheduled for later execution.

	Component	Action	Message
	<i>RM/JS decides to suspend job z using CPR</i>		
1	Autonomic Script	Throw	Predict Problem Node (node x)
	<i>Information gathered indicates emanate failure of node x</i>		
2	RM/JS	Catch	Predict Problem Node (node x)
	<i>Suspend scheduling on node x (predicted failure)</i>		
	<i>Translate node x to job z</i>		
3	RM/JS	Throw	Suspend Job (job z)
	<i>Suspend job z</i>		
4	MPI	Catch	Suspend Job (job z)
	<i>Coordinate a global checkpoint operation. Suspend/Terminate job z</i>		
5	MPI	Throw	Resume Job Cmd (job z)
	<i>Provide RM/JS with the command needed to resume job z</i>		
6	RM/JS	Catch	Resume Job Cmd (job z)
	<i>Store command with information for job z</i>		
	<i>Reschedule job z</i>		
	<i>Job z becomes runnable once again</i>		
7	RM/JS	Throw	Resume Job (job z)
	<i>Use stored resume information for job z to restart job</i>		
8	MPI	Catch	Resume Job (job z)
	<i>Bring job z back into a running state</i>		
	<i>Time passes, node x returned to service</i>		
9	Autonomic Script	Throw	Restored Node (node x)
	<i>Information gathered indicates node x is stable again</i>		
10	RM/JS	Catch	Restored Node (node x)
	<i>Return node x to resource pool</i>		

2.3 Workflow: Predicted Failure, Process Migration

A monitoring system predicts a node failure based on heuristic information gathered from the operating system, network card, and other system resources. The job is suspended and rescheduled for later execution.

	Component	Action	Message
	<i>RM/JS decides to suspend job z using CPR</i>		
1	Autonomic Script	Throw	Predict Problem Node (node x)
	<i>Information gathered indicates emanate failure of node x</i>		
2	RM/JS	Catch	Predict Problem Node (node x)
	<i>Suspend scheduling on node x (predicted failure)</i>		
	<i>Translate node x to job z</i>		
3	RM/JS	Throw	Migrate Node (job z: node x,q)
	<i>Allocate spare node q to job z</i>		
	<i>Migrate processes from job z on node x to new node q</i>		
4	MPI	Catch	Migrate Node (job z: node x,q)
	<i>Coordinate a global checkpoint operation.</i>		
	<i>Migrate ranks from node x to new node q. Resume application</i>		
5	MPI	Throw	Migrate Node Done (job z: node x,q)
	<i>Tell RM/JS that migration is finished</i>		
6	RM/JS	Catch	Migrate Node Done (job z: node x,q)
	<i>Receive confirmation that node x no longer contains MPI ranks</i>		
	<i>Time passes, node x returned to service</i>		
7	Autonomic Script	Throw	Restored Node (node x)
	<i>Information gathered indicates node x is stable again</i>		
8	RM/JS	Catch	Restored Node (node x)
	<i>Return node x to resource pool</i>		

3 Workflow: Faulty Interconnect

The following table details the events that each component will want to either throw or catch.

	Component	Action	Message
	<i>Initialization & Job Launch</i>		
0	RM/JS	Register	Failed Physical Interface (iface *: node *)
0	RM/JS	Register	Failed MPI Physical Interface (iface *: node *: job *)
0	RM/JS	Register	Restored Physical Interface (iface *: node *)
0	RM/JS	Register	Restored MPI Physical Interface (iface *: node *: job *)
0	RM/JS	Register	MPI Message Corruption (node *: job *)
0	IB Fault Monitor	Register	Failed Physical Interface (iface *: node *)
0	IB Fault Monitor	Register	Restored Physical Interface (iface *: node *)
0	IB Fault Monitor	Register	Check Physical Interface (iface *: node *)
0	Autonomic Script	Register	Failed Physical Interface (iface *: node *)
0	Autonomic Script	Register	Restored Physical Interface (iface *: node *)
0	Autonomic Script	Register	Check Physical Interface (iface *: node *)
0	MPI	Register	Failed MPI Physical Interface (iface *: node *: job z)
0	MPI	Register	Restored MPI Physical Interface (iface *: node *: job z)
0	MPI	Register	MPI Message Corruption (node *: job z)

3.1 Workflow: Fail-over to an Alternative Device

A physical network interface fails, MPI fails-over to an alternative device and continues.

	Component	Action	Message
	<i>Interface p fails on node x, job z running on node x</i>		
	<i>IB Fault Monitor is first to detect</i>		
1	IB Fault Monitor	Throw	Failed Physical Interface (iface p: node x)
	<i>Interface p on node x has failed</i>		
2(a)	RM/JS	Catch	Failed Physical Interface (iface p: node x)
	<i>Translate node x to job z</i>		
2(b)	Autonomic Script	Catch	Failed Physical Interface (iface p: node x)
	<i>Attempt diagnose and clean up IB routes and switches</i>		
3	RM/JS	Throw	Failed MPI Physical Interface (iface p: node x: job z)
	<i>Notify MPI of failed interface</i>		
4	MPI	Catch	Failed MPI Physical Interface (iface p: node x: job z)
	<i>Mark interface p as down</i>		
	<i>If possible, use an alternative interface</i>		
	<i>If not, suspend communication until interface restored</i>		
	<i>Interface p returned to service on node x</i>		
5	Autonomic Script	Throw	Restored Physical Interface (iface p: node x)
	<i>Interface p has been restored to service on node x</i>		
6(a)	IB Fault Monitor	Catch	Restored Physical Interface (iface p: node x)
	<i>Confirm interface is restored</i>		
6(b)	RM/JS	Catch	Restored Physical Interface (iface p: node x)
	<i>Translate node x to job z</i>		
7	RM/JS	Throw	Restored MPI Physical Interface (iface p: node x: job z)
	<i>Notify MPI of restored/new interface p</i>		
8	MPI	Catch	Restored MPI Physical Interface (iface p: node x: job z)
	<i>Add interface p back to the possible interfaces for communication</i>		

3.2 Workflow: React to Corrupted or Missing Data

A physical network interface is dropping or corrupting packets. MPI takes corrective action to mask such fails. At some point MPI may decide to remove the interface from service similar to Section 3.1

	Component	Action	Message
	<i>Interface p dropping or corrupting packets on node x</i> <i>MPI is first to detect</i>		
1	MPI	Throw	MPI Message Corruption (node x: job z) <i>MPI detects message corruption</i> <i>Continue masking corruption while interfaces are inspected</i>
2	RM/JS	Catch	MPI Message Corruption (node x: job z) <i>Translate node x to iface p-q</i>
3	RM/JS	Throw	Check Interface (iface p-q: node x) <i>Ask script to check interfaces for suspected failure</i>
4(a)	Autonomic Script	Catch	Check Interface (iface p-q: node x) <i>Checks interfaces</i>
4(b)	IB Fault Monitor	Catch	Check Interface (iface p-q: node x) <i>Checks interfaces</i>
5	Autonomic Script	Throw	Failed Physical Interface (iface p: node x) <i>Notify of confirmed failed interface</i>
6	RM/JS	Catch	Failed Physical Interface (iface p: node x) <i>Translate node x to job z</i>
7	RM/JS	Throw	Failed MPI Physical Interface (iface p: node x: job z) <i>Notify MPI of failed interface</i>
8	MPI	Catch	Failed MPI Physical Interface (iface p: node x: job z) <i>Mark interface p as down</i> <i>If possible, use an alternative interface</i> <i>If not, suspend communication until interface restored</i>
	<i>Interface p returned to service on node x</i>		
9	Autonomic Script	Throw	Restored Physical Interface (iface p: node x) <i>Interface p has been restored to service on node x</i>
10(a)	IB Fault Monitor	Catch	Restored Physical Interface (iface p: node x) <i>Confirm interface is restored</i>
10(b)	RM/JS	Catch	Restored Physical Interface (iface p: node x) <i>Translate node x to job z</i>
11	RM/JS	Throw	Restored MPI Physical Interface (iface p: node x: job z) <i>Notify MPI of restored/new interface p</i>
12	MPI	Catch	Restored MPI Physical Interface (iface p: node x: job z) <i>Add interface p back to the possible interfaces for communication</i>